

Defendant's Exhibit 20



September 10, 2021

VIA ELECTRONIC MAIL

Martie Kutscher Clark
Gibson Dunn & Crutcher LLP
1881 Page Mill Road
Palo Alto, CA 94304-1211
mkutscherclark@gibsondunn.com

Russell H. Falconer
Gibson Dunn & Crutcher LLP
2001 Ross Avenue, Suite 2100
Dallas, TX 75201
rfalconer@gibsondunn.com

Deborah L. Stein
Gibson Dunn & Crutcher LLP
333 South Grand Avenue
Los Angeles, CA 90071-3197
dstein@gibsondunn.com

Colin B. Davis
Gibson Dunn & Crutcher LLP
3161 Michelson Drive,
Irvine, CA 92612-4412 USA
cdavis@gibsondunn.com

Laura C. Mumm
Gibson Dunn & Crutcher LLP
200 Park Avenue
New York, NY 10166-0193
lmumm@gibsondunn.com

Re: *In re Facebook, Inc. Consumer Privacy User Profile*,
Northern District of California Case No. 3:18-md-02843-VC

Dear Counsel:

At the mediators' request, we send this message as a final attempt to avoid impasse on Facebook's production of the named plaintiffs' data. We look forward to your response on September 16.

At its core, this case is about what "content and information" (Facebook's term for data and information as set forth in its own terms of service) Facebook took from Plaintiffs, what Facebook told Plaintiffs it would do with their content and information and what Facebook *actually* did with it. This includes, but is not limited to, sharing it with third parties, negligently allowing third parties to take it and use it for improper purposes and failing to monitor third parties' use, which is precisely what occurred with Cambridge Analytica and Dr. Kogan, the scandal that sparked these consolidated actions. Facebook asserts it disclosed all its practices as to users' content and information and that users consented to those practices. To test this assertion, Plaintiffs need to understand whether Facebook's actions matched the conduct Facebook describes in its terms of service and privacy policies.

Gibson Dunn & Crutcher LLP
September 10, 2021
Page 2

KELLER ROHRBACK L.L.P.
BLEICHMAR FONTI & AULD LLP

Plaintiffs opened discovery in this case with the modest request that Facebook describe and identify the kinds of data it has collected on only the nine Named Plaintiffs, as opposed to the hundreds of millions of class members in this action. We have simply asked: what did Facebook collect about users and what did it do with it? RFP No. 9 seeks all documents Facebook has relating to the Named Plaintiffs, including the content and information collected about each of them.. RFP No. 10 seeks documents sufficient to identify the categories of “content and information” Facebook collects, tracks, and maintains about each Named Plaintiff. It has been Plaintiffs’ hope that this modest request could serve as a road map for class-wide discovery.

In Discovery Order No. 9, Judge Corley agreed. She identified the proper scope of discovery related to the data Facebook accumulates about the Named Plaintiffs as: (1) data collected from a user’s on-platform activity; (2) data obtained from third parties regarding a user’s off-platform activity; and (3) data inferred from a user’s on- or off-platform activity. Dkt. No. 557.

To date, Facebook has not produced data from categories 2 or 3. Such a production would include Facebook’s profiles of the Named Plaintiffs, and data Facebook acquires through its agreements with business partners, including data it bought and sold about Named Plaintiffs from data brokers in the heart of the Class Period. The parties have conferred and communicated at length about this issue, both before and after Judge Corley’s order.

Instead, Facebook continues to limit discovery to category 1. Facebook has repeatedly told Plaintiffs it has “produced the information contained in the DYI file for each of the Named Plaintiffs, plus certain additional information (such as a spreadsheet containing data tracking how Plaintiffs adjusted their Facebook privacy settings).” *E.g.*, Apr. 1, 2021 letter at 2 (attached). The external tool that Facebook created to share with Plaintiffs some small subset of the information Facebook collects about them does not meet the scope of discovery Judge Corley identified or even address the heart of Plaintiffs’ claims in this case. Plaintiffs know what they shared on the platform. But Plaintiffs do not know what Facebook collects, infers, embeds and tracks to create data sets about the Plaintiffs. Plaintiffs want to see those data sets and they want to know how Facebook uses them. Plaintiffs can then compare those actions to Facebook’s disclosures and the parties can have a meaningful dialogue about the scope of consent.

Because of Facebook’s secrecy and refusal to be transparent, there is a significant information asymmetry. Thus, it is impossible for Plaintiffs to identify with specificity the full scope of information Facebook has not produced about the Named Plaintiffs. But some internal Facebook documents give a clue to the types of information it collects about users. For example, Dep. Ex. 3 (attached) defines three broad categories of data Facebook “receive[s] about people”: native data, appended data, and behavioral data. *See* Ex. 3 at FB-CA-MDL-00213424. For those types of data, Facebook identifies categories of data it explicitly collects,

Gibson Dunn & Crutcher LLP
 September 10, 2021
 Page 3

KELLER ROHRBACK L.L.P.
BLEICHMAR FONTI & AULD LLP

implicitly collects, and infers. It appears that the Named Plaintiffs' data Facebook has produced is limited to information in the explicit collection category: profile information; posts, likes, shares; and location (checkins). Facebook has *not* produced all of the Named Plaintiffs' data it implicitly collects—location, phone number, carrier, device type, device identifiers. And Facebook has *not* produced the Named Plaintiffs' data it infers—interests/behaviors, public records, auto registration data, supermarket loyalty cards, retail purchases, credit card purchases, existing customer relationships, purchase history, data from customized third-party data, data from “enhanced” customer databases, website browsing behavior, conversions off Facebook, explicit actions off Facebook, mobile apps installed, activity within apps, or all device network activity. Nor has Facebook disclosed the extent to which it shares or makes accessible some or all of this data to third parties and what it does to monitor third parties' use of it.

The document at Bates No. FB-CA-MDL-00178902 provides further insight into the types of information Facebook collects about its users that it has not produced (limited to the Named Plaintiffs) here. Summarizing the value proposition of being able to read data from Facebook's platform, Sam Lessin writes: “We, Facebook, have more information based on our own derivation, as well as aggregate information about the connections/relationships of a user which we will give you if your service meets certain requirements and you are willing to pay us.” Among the data Lessin says Facebook has about each user is “[a]ggregate data about the tastes, properties, etc. of a user's friends (things this user's friends like, places they live, etc.),” “[d]erived data about a user / facebook's data/opinion of a user (probable location, account trust score, account age, etc.),” and “[d]ata provided by third parties — information which third parties have contributed to the graph on behalf of a user.” While the context for these descriptions is Facebook's consideration of shifting its business model to one where third-party developers paid for access to its platform, the descriptions quoted above reflect the data Facebook actually collects on its users. Facebook has *not* produced aggregated data about the Named Plaintiffs' friends, derived data about the Named Plaintiffs, Facebook's opinions of the Named Plaintiffs, or data provided by third parties to the graph about the Named Plaintiffs.

Facebook's patents also lend a clue into the kinds of data collects. For example, Facebook holds a patent titled “Determining user Personality Characteristics From Social Network System Communications and Characteristics” (U.S. Patent No. 9740752), which can be used to identify personality characteristics (e.g., extroversion, agreeableness, conscientiousness, emotional stability, and openness), which can be targeted by marketers on Facebook. Another patent, “Receiving Information About a User from a Third Party Application Based On Action Types” describes how “linguistic data and non-linguistic data associated with the user” are used “in a trained model to predict one or more personality characteristics for the user.” These “inferred personality characteristics are stored in connection with the user's provide, and may be used for targeting, ranking, selecting versions of products, and various other purposes.” (U.S. Patent 8732802B2 at 2). Examples of personality characteristics include: “extroversion, agreeableness, conscientiousness, emotional stability, and

Gibson Dunn & Crutcher LLP
September 10, 2021
Page 4

KELLER ROHRBACK L.L.P.
BLEICHMAR FONTI & AULD LLP

openness.” The patent further explains that “[e]ach user of the social networking system is associated with a user profile, which is stored in the user profile store. A user profile includes declarative information about the user that was explicitly shared by the user, and may also include profile information inferred by the social networking system. In one embodiment, a user profile includes multiple data fields, each field describing one or more attributes of the corresponding user of the social networking systems.” Each of these patents identify the types of information about Facebook users, including the Named Plaintiffs, that Plaintiffs have sought but which Facebook has not yet produced. These references are not intended to be complete. Rather, they are only included for the purpose of providing examples of Facebook acknowledging the existence of some aspects of the data Plaintiffs seek.

If Facebook will not make a complete production of all data it has collected about these nine people, Facebook must identify what it is withholding and why. And even if Facebook does make a complete production, the parties must confer in advance of the production in order to agree on the form of those productions, given the complexities with the kinds of data collected, inferred, aggregated and used.

Regards,



Derek W. Loeser
dloeser@kellerrohrback.com



Lesley E. Weaver
lweaver@bfalaw.com

ATTACHMENT A

.

From: Simone LiTrenta </O=THEFACEBOOK/OU=EXCHANGE ADMINISTRATIVE GROUP (FYDIBOHF23SPDLT)/CN=RECIPIENTS/CN=SIMONEL051>
Sent: Thursday, May 08, 2014 4:34 PM
To: Matt Scutari; Rob Sherman; Emily Sharpe; Emily Vacher; Maritza Johnson; Travis Bright
Cc: Erin Egan
Subject: Offsite presentation
Attachments: Combined papers.docx

Hey, all. Attached is the combined doc of privacy team papers that Marne will be sending out with the rest of global policy team papers.

Now on to the slide deck☺. If someone has started slides already and can share the format with the group so we can make the look and feel uniform, that would be great. If someone is a PPT genius and wants to take the lead on combining finished slides into one deck and making minor changes, let me know. Otherwise, please save your slides to the folder for slides and I can combine.

<https://www.dropbox.com/sh/hquhjw021dr3qty/AADBwXtmeiNZJRBZR6-xjlgia>

Erin would like to review the slide on the plane Monday morning. If everyone can finish their slides by COB Friday/Saturday, I can get them to her Sunday night.

If you would like to discuss your slides with Erin, please let me know ASAP and I will find time manana.

Simone

Hard Questions

Does Facebook share my data with advertisers?

"We don't share the private information that you put on Facebook with advertisers without your consent."

Why do we use that framing?

- Though Facebook's policies prohibit sharing of data with data brokers or similar entities, we only make commitments about what Facebook will do.
- Advertisers who are also developers might get user information through Platform, with user consent.
- When people make information public, anyone can see it on or off Facebook, so we don't make any non-disclosure commitment for that information.
- If someone else shares information (including information about you) they control who sees that information. Also, we may in the future operate a "data cooperative" or other product allowing exchange of information that people haven't given to us directly.

How does ad targeting work?

Key types of targeting

- *Core demographic and interest targeting.*
- *Facebook Exchange.*
- *Custom audiences and partner categories*

How can people see what you know about them and control their ad experiences?

Context menu

- Click the "x" or "v" next to any Facebook ad to see who the advertiser is and hide that ad or other ads from the same company.

Activity Log and Download Your Information (DYI)

- See the information you've put on Facebook that may be used for ads. Delete or change who can see it.

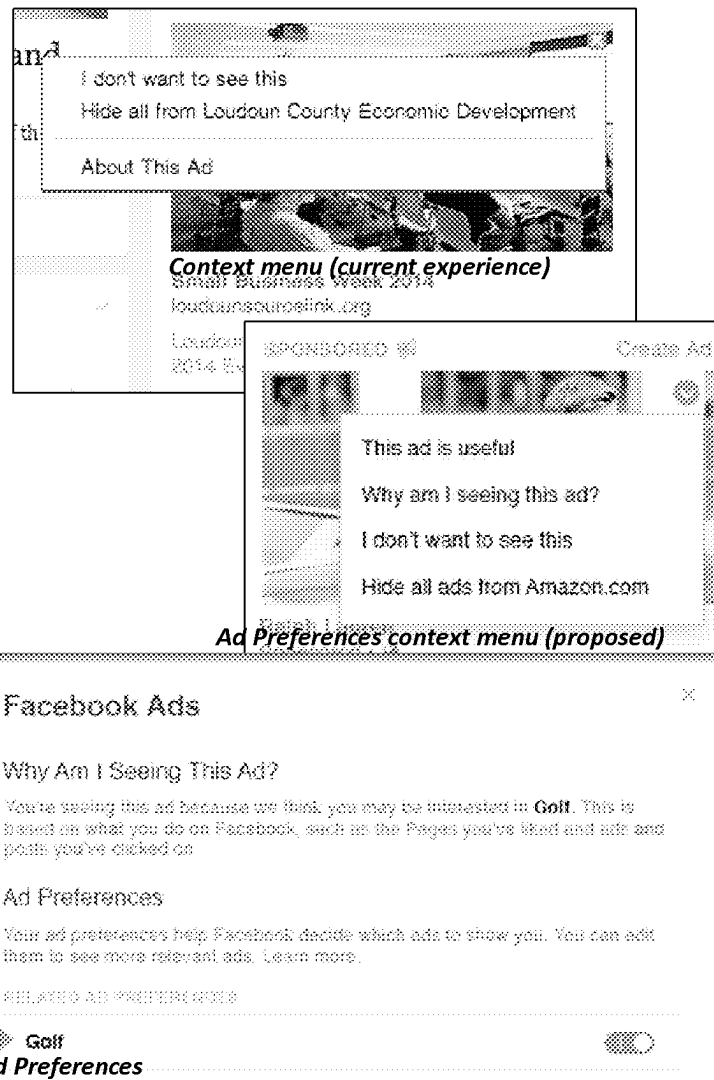
Ad Preferences (launching summer 2014)


- See why you saw a particular ad and what you can do about it. Also see and modify all your ad clusters.

Centralized opt-out

- Turn off collection and use of third-party behavioral data with one click – not just on FB but across the web.

See mocks on next page.



Behavioral Advertising Announcement*(Tentative plans subject to change; late May/early June)***Learn More:**

Hi Jane – we want to explain the choices you have over the ads you see.

- ***You can choose to opt out:*** We want to show you more meaningful ads based on what you do on and off of Facebook, but you can opt out of how this is done by visiting the [Digital Advertising Alliance \(DAA\)](#).
- ***You can choose ad categories:*** Today we're announcing a new tool called Ad Preferences, which lets you choose info that will influence the ads you see. It's launching today in the US, and will be available everywhere soon.

Read more about:

- [Ad Preferences](#)
- [How we use cookies, pixels, and similar technologies to show you ads on and off Facebook](#)

Location

(Maritza Johnson)

Knowing where people are when they interact with our services is useful for designing innovative products, customizing content, and improving our security features. However, location data is generally considered to be sensitive data that deserves a greater level of attention.

In general, we consider the following questions when new products are proposed that leverage location:

- When is the location data collected? Is location data collected when the app is in the foreground (the app is the primary content on the screen and the user is actively engaged), or in the background (the app is running on the phone but the user is not actively engaged with the app)?
- Frequency?
- How is the data point collected? It's possible to collect the device's location with varying degrees of accuracy (GPS, WiFi, Bluetooth, IP addresses, MAC addresses, etc.).
- What level of granularity is the location data collected and stored? How accurate is the data point?
- Retention, how long do we store it?
- What controls are available?
- Is derivative data produced (ex., ad clusters)?

Additional considerations depending on the nature of the product and the use of location data

In cases like Nearby Friends (Aura) where people are intentionally interacting with the feature because it is location based, we offer additional controls and review options so that the user understands what is collected. The combination of the controls, the transparency offered by "Travel History," and the opt-in nature of the product mitigate privacy concerns.

We also offer products that rely on location information as one of many inputs, in some cases, the data used to represent location is less granular or could potentially be less accurate. In those cases, we take a close look at the use of the data, the granularity of the data, the retention policy, and how the data will be stored to minimize the chance that individuals could be linked to the data collected.

Hard Questions

Are you tracking where I go all the time?

No, unless you recently turned on "Nearby Friends." Visit your Travel History to see what Facebook knows about your recent whereabouts. If you have chosen not to use this feature, then the Facebook app will only collect your precise location when you open the app if you allow the app to access your location.

How does Facebook use my location?

We offer several products that allow you to share your location with your friends like Nearby Friends and Places. We also use location data to customize content in a language that matches the general location where you are using the app and to offer security features.

When does Facebook collect my location?

Before the launch of Nearby Friends, the Facebook app did not collect location information when people left the app. The Facebook app collected location data each time the user launched the app, or when the user clicked on the composer box to start a new status.

How is my location used for advertising?

Facebook currently offers advertisers the ability to target ads based on where people live. For example, a restaurant in Kansas City could choose to show ads to people who live within 10 miles of the zip code where the restaurant is located. The location information Facebook currently uses to target ads comes from people's self-reported Current City, as well as other signals such as their IP address.

Are you using data from Nearby Friends or Location History to target ads?

No, at this time we are not using data from Nearby Friends or Location History to target ads. The launch of this product doesn't impact the way advertisers can target people based on location.

I want to hide my location from Facebook, how can I turn that off?

If you're using a mobile app, go to the Settings section on your device and turn off location for the Facebook app. Your IP address will still be used to approximate your location for security features and content localization.

Can law enforcement get my location history now?

Your location history is treated in the same way as other data that Facebook collects about you. Law enforcement officials need to go through the same processes to request it. You can also always delete your location history from the Location History section of your Activity Log.

General Context

- Location-aware services and advertising have become increasingly popular and feasible with the growing popularity of smartphones and other personal devices.
- Geolocation data is recognized as one way that the online and offline worlds can be linked which is a great feature for products, however it introduces additional privacy concerns.
- Regulators worldwide are paying close attention to location data and have said that the sensitive data deserves a greater level of privacy protection.

Products that use location:

- Nearby Friends (Aura), <https://www.facebook.com/help/629537553762715/>
- Travel History
- Site Integrity features (identifying suspicious logins by location of attempts)
- Hunch (to offer local recommendations based on current location)
- FB WiFi (opt-in)
- Measurement for ads – offline measurement, Clicker, Demographics of the World (utilizing the MAC address when people enter a specific range).
- Ads targeting

Additional background on the uniqueness of location data

In a study released in March, researchers analyzed the coarse location histories recorded to the nearest hour of 1.5 million mobile phone users. The researchers found that knowing which 23-block area a device was in during four distinct 60-minute periods was enough to uniquely identify 95 percent of location histories in the dataset.

<https://www.privacyassociation.org/resource-center/unique-in-the-crowd-the-privacy-bounds-of-human-mobility>

Facial Recognition

(Emily Sharpe)

Background: The facial recognition feature has been the subject of controversy since its launch. After we acquired Face.com, which previously provided the software behind tag suggestions, we took the product offline in order to complete performance improvements. We re-launched the product in the Spring of 2013 in all markets except Canada and Europe.

Current status of Facebook's facial recognition practices: Tag suggestions is an opt-out feature that uses facial recognition technology to help suggest friends that people might want to tag in photos. Unlike other services that offer facial recognition, we only suggest you to people you're already friends with.

- **Notice and Control:** You're informed when someone uploads a photo you're tagged in, and you get added to the audience so you can see it. You can choose to remove the tag, ask the person to remove the photo, or report it to Facebook. Tag review allows you to approve a tag anytime someone tags a photo you're in. You can also turn off tag suggestions so Facebook won't suggest that people tag you when photos look like you, and the template that FB created to enable the tag suggestions feature will also be deleted (although friends will still be able to tag photos of you).
- **Security:** Facial recognition templates are stored in encrypted form, and are built in a format that is intentionally not interoperable with other systems so even if someone were able to access them and decrypt them they would not be useful to a third party. We also implement legal and technical measures to deter data harvesting.
- **Government access:** We apply strict legal and privacy requirements to all law enforcement requests for photos and facial recognition templates.

Policy considerations:

- Regulators in Canada and Europe have expressed strong concerns about our use of facial recognition technology for people who live in their jurisdictions, and the Irish regulator particularly threatened to give us an unfavorable audit report if we didn't commit to stop using facial recognition in Europe without an "opt-in" by each use.
- In the U.S., the National Telecommunications & Information Administration (NTIA) is facilitating a privacy multistakeholder process focused on developing a voluntary, enforceable code of conduct that specifies how the White House's Consumer Privacy Bill of Rights applies to facial recognition. Facebook's facial recognition practices and research have garnered a fair amount of attention during the process, and FB and Google have been pressured by one privacy advocate to present on our practices. Some participants have also focused on FB's large photo collection as being at risk of data harvesting. While we believe our current and planned practices generally align with consensus policy positions, we are working with industry associations to shape the code of conduct.

Our commitments:

- We've committed that we won't turn the feature on in Europe unless we agree with the Irish DPC on the type of consent we'll put in place, and that we will consult with the Canadian Privacy Commissioner's Office before offering facial recognition in their country.
- We would likely support (or not actively oppose) NTIA code of conduct principles articulating that 1) companies should tell individuals that they are collecting their facial recognition data, and 2) companies should provide subjects of images with meaningful control over how their images are used to share information with others who would not otherwise know that information.

Product updates: In 2014 we're working on a few changes to the way we use facial recognition technology.

- **Face Alerts:** A new, opt-in feature that allows people to be notified about photos they're in on Facebook, but haven't been tagged in. They will only be notified about photos they're in the audience to see.

- Expected launch is early July (based on EU's team's feedback), and will be available worldwide, including the EU and Canada.
- When people opt in to Face Alerts, they'll see these photos in a new section on the "photos of you" page, and these photos will only be available to them.
- People can choose to tag these photos to make it part of their identity and timeline. They can also reject the tag suggestion, preventing people from tagging them in that photo; ask the person who uploaded it to remove the photo; or report it to Facebook.
- This feature will be available if your "tag suggestions" setting is turned on.
- **Graph Search.** We're looking to incorporate facial recognition results into Graph Search. For celebrities, we would use facial recognition to automatically include photos of a celebrity in search results. For non-celebrities, we wouldn't automatically integrate photos into search results but would display a distinct "Do you want to tag Sheryl in these photos?" unit next to search results that are informed by facial recognition data.
- **Storing face templates:** The Privacy XFN team is currently exploring the feasibility of storing face templates in both the US and EU even after a user turns off the tag suggest setting. The purpose would be to continue using facial recognition for internal suggestions/security features, without suggesting photos for friends to tag. We would offer users the option to delete their templates when they turn off the setting, either right in the flow or by linking to the Help Center and providing a contact form.
- **Videos tag suggestions:** We're rolling out a feature that suggest tags in videos. This feature works just like tag suggestions in photos and, for now, is only available on the web when the uploader of the video clicks the "edit" button. The underlying technology and settings/controls are the same as tag suggest for photos. As with all tag suggest features, this will only work if people in the video are your friends. We are using reactive messaging here in case concerns about "real time facial recognition" surface.

Hard Questions:

Do you support regulation of facial recognition technology?

People's greatest concerns about facial recognition technology don't generally arise from problems with the underlying technology itself, but to problems with how certain people might *use* photos in a way that is inconsistent with individuals' expectations. We don't think it's helpful to regulate the use of specific technologies. This only encourages bad actors to come up with new ways to do the same things that aren't yet regulated. A better approach is to focus on specific activities that raise concern, regardless of what technology is used for those activities. We support technology neutrality – making sure that regulations focus on practices that raise concerns, not technologies that have both positive and negative uses.

Facebook talks a lot about user control and its privacy protections, but doesn't that all go out the window when the government gets a warrant and collects all of that data?

Our facial recognition technology is proprietary and therefore not particularly useful to government investigators. Our facial recognition software generates a template, which is basically an average of the characteristics of multiple photos of a person. That template is encrypted and requires a "key" to decode it. It can't reliably be used to recognize an individual from among thousands or millions of templates, so if a law enforcement agency tried to use our templates in that way, it wouldn't work properly. We also don't have any reason to believe that existing laws would allow government agencies to compel Facebook to cooperate in a "fishing expedition" for photos, and, if they did, we would vigorously defend the privacy of our users. We aggressively protect our users' data when confronted with law enforcement requests: We scrutinize every government data request that we receive – whether from state, local, federal, or foreign governments. We frequently reject such requests outright, or require the government to substantially scale down its requests, or simply give the government much less data than it has requested. And we respond only as required by law.

Facebook has the world's largest database of photos. What are you doing to prevent people from "scraping" the photos so they can combine them with publicly available FB profile data, and then apply commercially available facial recognition technology to identify otherwise-anonymous people walking on the streets or in other photos online?

We implement legal and technical measures to deter data harvesting (a.k.a. "scraping"). Although we do not discuss specific technical measures for security reasons, we use rate limiters and other methods to detect and deter efforts to harvest publicly available Facebook data. Our policies also prohibit people who access our site from screen-scraping and similar data harvesting efforts. We have taken legal action where we have learned that people have attempted to circumvent these measures.

Apps, Acquisition, and Creative Labs

(Travis Bright)

For the last few years Facebook has moved towards a federated system of applications compared to our previous work to integrate all featured into the main Facebook app. In addition, we acquired Instagram that has continued to run pretty independent from the rest of Facebook and Mark's statements around WhatsApp indicate that they will run all but completely separate from Facebook. With Creative Lab's focus on incubating new ideas in the form of small projects brought to market quickly (and often as a stand alone app) which increases the trend of "ever more apps". There are some fundamental questions that have to be addressed:

- Do people use the system anonymously, pseudo-anonymously, or with their real name?
- What level of data integration will exist? Nothing, anonymized/aggregated/hashed data mapping, or full data integration (increasing the person's Facebook profile)
- What level of separation are we trying to achieve from the main brand?

Level of Anonymity

Facebook has long championed our real name culture. But more and more of our competitors and our own internal apps are breaking away from this because people often perceive real name system with being "heavy weight". To make a fun, lighter experience that often is more ephemeral than traditional Facebook our PMs and Designers are moving towards no-name or aliases. This complicates our external messaging and our ability to react to abuses but what is critically important is our company's ability to test new things. Our product teams have to be free to push the boundaries and create new test without being constrained by our previous positions and thoughts. As a policy team that has to speak externally, build relationships, and judge the impact on external people each change will bring this is and will continue to be a problem. We'll need to make sure that all of our statements are focused on the current system and always leave the expectation that change will happen.

Data Integration

Facebook's data is hugely valuable but comes with a lot of restrictions we've either placed on ourselves or by external parties (regulators). Some apps want to take advantage of the data we have while some are trying to simplify their app by running it independently. For example, Slingshot's data is stored in Parse completely separate from Facebook data. You only need a phone to create an account, aliases used in the app aren't linked to Facebook profiles, and they are showing ads so don't even need demographic or aggregated data. But even Slingshot needed APIs built to allow Facebook to extract data for reporting and search warrants. They built an independent contact importer which only stores other Slingshot user's (no grey accounts). They will offer an option to add your Facebook friends using existing app SDKs that will be the first bit of integration. At that point Facebook would know that a particular Facebook user also used Slingshot as well as Slingshot knowing their user is also a Facebook user.

The next step up from this is sharing of anonymized, aggregated, or hashed data. This allows two systems (normally Facebook and one of the apps) to map their own data in different ways without actually mixing the data. An example of this is Atlas that needs Facebook to de-duplicate people across browsers and devices. Traditionally ad data and conversion tracking relied on cookies. A cookie would be written when Bob saw an ad and then if he later made a purchase the conversion tracking pixel would read the cookie and match the two actions together. The problem here is that each browser on each device has its own cookie storage (and mobile apps don't have the same cookie mechanism). So if Bob sees an ad within a Firefox window and then later makes a purchase within a Chrome window or via his iPhone the ad server and measurement would not be able to link those actions together leading to poor knowledge about what's driving people their purchases. Facebook can solve this by writing a hashed UID to the Atlas cookie. Since your Facebook UID is the same across all of your browsers and devices you are logged into this means that the hashed UID will be the same on all of these systems. Atlas can then use that unique identifier to link different actions across browsers and devices.

Finally there is data integration. This can range from a small amount of data to full integration. But I'm defining this as where data on one or both sides (Facebook and the app) gain data and merge that with their own data. An example of this is the Messenger app which stores its data separately but has access to your full Facebook profile, adds to it (messages and changing coefficient as you use it). Arguably a lighter version of this is any app that uses Facebook login since Facebook gains knowledge of your use of the app and adds this to your profile. Paper doesn't really fall into this bucket because their data is Facebook's data. The app is basically just a different skin/frontend to Facebook.

One Brand or Many?

A more recent trend is for product teams to want to distance themselves from Facebook by not directly referencing it or obscuring the link. They feel that different groups of people feel are more concerned with sharing certain data when they think about Facebook. They believe that people mentally associate their friends, co-workers, family members, ... along with permanence with Facebook. They don't want Facebook in the name, don't want Facebook in the privacy policy, shown during registration, ... And this isn't just for Facebook. Pepper, the new direct sharing app for Instagram is likely to have an external name that doesn't use "Instagram" in it. We obviously can't hide our ownership of these apps but getting back to the product team's need to be able to test things we need to allow them some flexibility to break away. This also comes up when we acquire an app. Obviously, Moves and WhatsApp have never mentioned Facebook in their product before and the mindset seems around keeping them as separate as possible going forward. They have different privacy policies, different content policies, ... From a tactical side we need to figure out a way to deal with a world where we have dozens of different DUPs, Privacy Policies, and rules. We need to decide for our internally built apps whether we will allow them the all to have their own custom policies or if we should build out a few tiers of policies and ask them to select the level they need.

Hard Questions:

Where do we start with a new app or acquisition? The basic questions:

- Data Policy
 - What data is collected? (PII, location, ...)
 - Where does is the data stored?
 - How long is the data stored?
 - Is there any sharing of data between apps and/or Facebook?
 - Will we target ads with this data?
 - Will this system use Facebook's DUP?
- Content Policy
 - What are the rules for using this new system and are they different than Facebook's?
- People Policy
 - Will the system know/collect ages?
 - What is the minimum age?
 - Is this system using real names?
 - Is there reporting? Blocking?
 - Does the system verify accounts? How?
 - Will there be ads? From where?
 - Will we charge for this service?
- Privacy Policy
 - What is the default level of privacy for the system?
 - If this is an existing system what previous promises were made?
 - How are we going to fulfill our Irish DPA/FTC obligations for this system?
 - Will this system use Facebook's Privacy Policy?

There is a quick matrix of common apps and their differences here:

<https://docs.fb.com/sheet/ropen.do?rid=osbge49eb2cbdfc004a8ea79e35865b97817b>

Safety Policy Paper **(Emily Vacher)**

What are our safety goals and priorities for 2014?

- (1) Public Policy.** It is our responsibility to demonstrate our position in the industry as thought leaders regarding safety issues and to support policy makers when they are faced with pressure to legislate in areas that are better handled by empowering users with tools and education.
- Review of proposed safety legislation/policies.
 - Safety is an area in which we typically do not want to encourage legislation or the appointment of eSafety officials. We will work with country managers/legal when safety related legislation is proposed to minimize the impact to FB.
 - Support in this area includes targeted education and resources for teens and adults, strong partnerships with key safety groups and organizations, and the development of technical solutions and effective reporting channels.
 - Identify highest risk jurisdictions and target with appropriate proactive educational campaigns, relationship building, and other strategic engagement.
 - At present, resources are being primarily directed to a few of the highest risk countries, including Australia, New Zealand, Canada, Brazil, US, Singapore (when the comprehensive review is completed, this will be reevaluated.)
- (2) External relationships and partnerships.**
- Develop and leverage strategically selected domestic and international child safety/advocacy partners who can provide feedback on programs and products, speak to policymakers and press as needed and demonstrate to policymakers and we are aligned with experts.
 - Partners include NCMEC, SAB, and international child safety organizations like IWF, SaferNet, Chicos.net, Project RockIt. We also work cross functionally with both Policy Comms and Security Comms to proactively present our messaging and to reactively deal with media fires and/or misinformation.
- (3) Product safety**
- Product Reviews: Review all upcoming product and feature releases to ensure they don't violate our core safety principles or put our external relationships at risk.
 - Work with cross-industry partners to drive safety innovation in technology, including in the areas of suicide prevention, social resolution, and PhotoDNA. Additionally, help NCMEC achieve their goals of finding missing kids and ending exploitation through tools like AMBER Alert v2 and video hashing.
- (4) Teen education and empowerment.**
- Build, pilot and scale a "Youth Ambassador Program": We are currently negotiating with ChildNet International to develop the curriculum for this program, and will initially pilot the program in London, the US, and one other market.
 - Work with Policy Programs to identify gaps and develop appropriate educational programs/materials to demonstrate our thought leadership.
 - Bullying Prevention Hub Global Rollout (EMEA, LatAm)
 - Think Before You Share Global Rollout
 - Friend in Need College Campus Tour/Town Halls
 - Internet 101 for Emerging Markets (in development)
 - Use data driven research to help us make decisions in safety policy, program development and for talking points with policymakers.

- Investigate research possibilities with FOSI, University of New Hampshire CAC Research Center, danah boyd, etc.
- Support country managers by providing safety materials they can use in their countries

(5) Adult education and empowerment.

- Engage with organizations who work with adults who parent, support, educate, and mentor children and provide tools, resources, and educational programs, including religious communities, teachers, Scouts, Big Brother/Big Sister, Cal Ripken, Sr. Foundation, pediatricians, school resource officers, etc.
 - FB 101 For Parents & “Having the Tech Talk”
 - Leverage FOSI’s new “GDP” (Good Digital Parenting) program for opportunities for co-branded parent focused FB educational materials/programs, distribution channels, and research opportunities.

Hard Questions

Why are you now letting kids 13-17 now post publicly? Isn’t this dangerous for them?

On Facebook, you control who you share with. That can be a single person in a message, a small group, with friends, or with the world. Before we updated our policies, for people aged 13 through 17, the initial audience of their first post on Facebook was set to “Friends of Friends” – with the option to change it. After this update, a new teen on Facebook’s first post is set to a narrow initial audience of “Friends” only. It’s important to realize that teens are among the savviest people using social media, and whether it comes to civic engagement, activism, or their thoughts on a new movie, they want to be heard. So with this update, people aged 13 through 17 now have the choice to post publicly on Facebook. While only a small fraction of teens using Facebook might choose to post publicly, this change now gives them the choice to share more broadly, just like on other social media services.

You say your policy is that you have to be 13 or over to have a FB account but studies have shown that there are millions of kids under 13 on Facebook. Why doesn’t Facebook do more to keep kids off the site?

- Safety experts agree that there is no singular way to keep underage children off of any platform – so we employ a layered approach to keeping kids under 13 off our site:
 - People who sign up for a Facebook account are required to type in their age on the very first screen. We use a simple piece of technology called a cookie to ensure that if a teen tries to sign up for FB and indicates that they are 12, they can’t simply come back a minute later and say that they’re 13+. This helps keep underage kids off of our site, and it’s something we’ve chosen to do (not required of us).
 - We ask people to notify us if they believe there is an underage user on our site; we have a dedicated compliance channel for these reports; and we delete the accounts of children under 13 as soon as we become aware of them.
- However, one report highlights that most parents with under 13’s on Facebook not only know their kids are on our service, but helped them sign up (and lie).
 - This report (by Danah Boyd, Eszter Hargittai, Jason Schultz and John Palfrey) notes that parents actively assist their children under 13 in joining Facebook even though they know it violates our policy.
 - The report also highlights the difficulty in implementing age restrictions on the Internet and underlines the need to continually work to keep kids safe online.
 - We appreciate the attention that these types of reports are giving to child safety, and believe that they provide important opportunities for everyone – industry, government, parents, advocates – to discuss it with the ultimate goal of trying to provide better, safer experiences for kids online.

What is Facebook doing to prevent bullying on its site?

Bullying in any form – from the school playground to online – is unacceptable. Facebook is the industry leader when it comes to bullying prevention, and we help people speak up for each other. We've created policies, programs, and tools to foster safety, accountability and trust in our community.

- **Policies:**

- *Real Name Culture* – We've built a network based on authentic identity, so people are more likely to treat each other with respect. It is a violation of our policies to use a fake name or operate under a false identity.
- *Easy to Use Reporting* – There are 'Report' links on virtually every Facebook page, and anyone in our community can block people who post hurtful content. Harassment, bullying and other forms of abuse are prioritized by our response team.
- *A Safe Experience for Minors* – Facebook's privacy and visibility settings take into account the unique needs of people between the ages 13 and 17, and are generally more restrictive than the settings for adults.

- **Product Innovation:**

- *Fostering Compassion* - We've partnered with the Yale Center for Emotional Intelligence to provide youth and adults with valuable tools and strategies to help them effectively address bullying behavior and its consequences.
- *Conflict Resolution* – We believe in the power of friends to help prevent bullying. That's why we've created tools like Social Reporting that enable people to report bullying and harassment to trusted friends and adults. Our data shows that in over 80% of cases, when a teen asks another teen to remove a photo from our site they don't like, the other teen will comply. Approximately 3.9 million people use this tool each week.
- *Support Dashboard* – Everyone should be able to know what happens to reports sent to Facebook. We've created a feature that lets people see whether or not their report has been reviewed and allow them to be notified when a decision is made (note: the dashboard does not cover every report made by users – e.g., if you report someone for using a fake name.)

- **Digital Citizenship Education:** Consumer education is a key pillar of creating a culture where people treat each other with respect. Together with safety advocates, academics, government and industry, we use our platform to raise awareness of bullying and how to prevent it.

- *Family Safety Center* – In November 2013 we launched an expansion of our Family Safety Center. We worked with the Yale Center for Emotional Intelligence to develop new content on bullying prevention, including tips, tools, and actual scripts for all stakeholders: parents, educators, and teens, even people accused of bullying. The new hub is also directly tied to reporting on Facebook because we know that it's so important to give people who get bullied or see it happening the resources they need at the moment they need it most.
- *Bullying Prevention Education* – Facebook has launched several education efforts aimed at inspiring bystanders to take action to stop bullying. In the US, we've partnered with Time Warner to create a social media pledge (Stop Bullying; Speak Up) to speak up when adults and students see bullying occur – today more than 140,000 people have taken the pledge.

Are you targeting children with ads that are not appropriate for them?

We have strict guidelines over what organizations can advertise to people 13 to 17 on Facebook. This is one reason it is important for teens who set up an account use their real age – to make sure they receive the many protections we offer teens in our community.

Contact Import

(Matt Scutari)

A Brief History of Facebook Contact Import

Facebook's contact import functionality (or "Find Friends") is designed to drive engagement and growth by helping people find and add new friends on Facebook. This is a three-step process:

1. Person uploads his or her contacts to Facebook.
2. Facebook shows the uploader contacts who are also on Facebook and prompts the uploader to add them as friends (drives engagement).
3. Facebook shows the uploader contacts who are not on Facebook and prompts the uploader to invite them to join Facebook (drives growth).

Over the years, contact import functionality has been the subject of significant legal and regulatory attention in various jurisdictions around the world (for Facebook and others). Facebook has mitigated risks associated with its Find Friends feature in a number of ways, including by displaying notice before a person uploads contacts, providing a way for people to manage and delete the contacts they upload, providing people with granular control over invitations, and restricting what types of data can be uploaded to Facebook's servers from the person's device.

Much of the controversy surrounding Facebook's Find Friends feature has involved our collection and use of the non-user data in connection with facilitating invitations to join Facebook. Particularly in Europe, non-users are often viewed as owning their personal data, regardless of the source. However, we take the position that any contact information uploaded by a Facebook user belongs to the uploader. While this argument has been challenged, it appears to be sustainable to the extent that we're using non-user data only to help the uploader connect with them on Facebook (and potentially to help others connect on Facebook).

Our collection and use of non-user data may become more problematic to the extent we attempt to use non-user data for other purposes. The Irish DPC's 2011 audit report addresses complaints regarding our alleged use of contact information uploaded by users to create so-called "shadow profiles" of non-users without their knowledge. We maintained that we do not create such "shadow profiles," and stated that non-user data is only used to allow the uploader to send invitations to non-users. Importantly, we relied on the argument that a non-user is clearly informed that Facebook has her contact information when someone sends her an invitation, which offers a link to allow non-users to delete their contact information (although note that we retain a non-usable hashed version of the non-user's contact information in order to prevent any further invitations from being sent to that address). The DPC ultimately found that "the upload of contacts by individuals to facilitate the sending of invitations to friends could operate in compliance with the Data Protection Acts provided full information was provided to non-users in relation to the use of their email address data on receipt of an invitation and any requests for removal are respected."

Other Facebook apps, such as Messenger, Slingshot, and Instagram also have contact import features, but they differ from Facebook's functionality in various ways. For example, Messenger, Slingshot, and Instagram all currently sync contacts on a periodic or "continuous" basis, while Facebook's contact syncing is still a manual process that must be initiated by the user each time (although this may change in the coming months). Additionally, contact import is required as a condition of using Slingshot, while it is optional for Facebook, Messenger, and Instagram. However, Slingshot does not retain any non-user data following a contact sync.

Note that WhatsApp has also been subject to significant regulatory scrutiny for its contact import functionality, particularly with respect to its collection and use of non-user data.

Hard Questions

Q. Does Facebook collect non-user data?

A. If people on Facebook choose to upload information about their contacts, Facebook will maintain this information on the uploader's behalf and use it to let the uploader invite people to join Facebook. We give non-users an opportunity to request that Facebook stop using their contact info for friend suggestions whenever a Facebook user sends them an invite to join Facebook. We keep a hashed version of their info to ensure that we honor such requests.

Q. What does Facebook do with non-user data?

A. We currently maintain contact information regarding non-users only to allow the Facebook user who uploaded the contacts to invite the non-users to join Facebook. [Note: The growth team is considering using this data for PYMK, so we should be careful to speak about this only in terms of what we do currently.]

Q. How does Facebook justify collecting and using non-user data when the people to whom the data pertains are not subject to Facebook's privacy policy?

A. At this time, we only allow people who uploaded their contacts to Facebook to send invitations to those contacts who are not on Facebook. If the uploader sends an invitation to a non-user, the non-user can request that Facebook stop sending invitations or using the person's contact info for suggestions.

Q. Does Facebook create "shadow profiles" for non-users?

A. No.

Q. What does it mean to "continuously sync" contacts?

A. "Continuous sync" is an industry buzzword that can mean a number of different things. Typically, "continuous sync" functionality is not literally continuous, but rather periodic.

Q. Does Facebook continuously sync contacts?

Facebook does not continuously sync contacts at this time, but we are always testing new features to make it easier for people to connect on Facebook. Although Messenger, Slingshot, and Instagram do continuously sync contacts, none of these processes are truly continuous. For example, Slingshot will sync with your address book each time you do a cold restart of the app, and Messenger will periodically check to see if you had added new contacts to your address book.

ATTACHMENT B

.

From: Douglas Purdy </O=THEFACEBOOK/OU=EXCHANGE ADMINISTRATIVE GROUP (FYDIBOHF23SPDLT)/CN=RECIPIENTS/CN=DOUGLAS PURDY>
Sent: Thursday, August 30, 2012 10:17 AM
To: Sam Lessin
Cc: Mike Vernal
Subject: Re: Platform Business Model Framing

I am on vacation this week (still), but happy to help.

That said, in wonder if you should just take the pen on the deck moving forward so I can have a day off and not be the bottle neck?

I can send you what I have so far in an hour when I get back to a computer.

On Aug 30, 2012, at 10:11 AM, "Sam Lessin" <sl@fb.com> wrote:

1. Sorry I wasn't clearer about my plans / being out. That is my bad guys / didn't mean to leave this hanging (obviously super important)
2. Doug, really glad you are framing this up / working on this ...
3. Here are my notes on where things are from Friday, I think a lot of it is covered below but I just want to make sure we are in sync here, are presenting all the hard questions for this offsite

(A) How the platform / our APIs would work if we could start over / the base rules for all businesses using Facebook.

- There are two basic sides to platform that function differently
 - There is a 'write' side, whereby businesses can get permission to write to the graph on behalf of users
 - There is a 'read' side, whereby businesses can get permission to read from the graph on behalf of users
- On the 'write' side:
 - The value proposition: you can write to our system both messages 'on behalf' of users (think explicit posts and timeline boxes), and messages on your own behalf (think page posts), to drive growth and re-engagement. We give you a natural amount of distribution for free / to make our user experience best, and we charge for everything else
 - Applications can write whatever they want to the graph on behalf of users (with the user's permission in most cases)
 - Distribution
 - All content written by applications gets the natural amount of newsfeed distribution based on the NF algorithms for maximizing engagement & user happiness
 - *Any / All applications can pay to up-rank themselves in feed if they want more traffic*

- *Any / All applications can pay to get into premium channels (invites, inbox, etc)*
 - Information
 - If you write structured data about a user in a way that other businesses want to use for targeting *we will pay you* (mechanism TBD) and you can set certain limitations on how the data can be used / by what parties.
- On the 'read' side:
 - The value proposition: you can read from our system. Users can always give you the information they have given us directly in order to help you customize your service / provide better service. We, Facebook, have more information based on our own derivation, as well as aggregate information about the connections/relationships of a user which we will give you if your service meets certain requirements and you are willing to pay us
 - Applications can use Facebook for 'login' to allow users of their application to not need to remember another password, etc.
 - User-Data
 - Applications can request a user give them 'their data' in order to provide the user with a more custom experience (*NB: the 'user's data' is explicitly the data which the user has entered themselves into Facebook - e.g. Name, profile photo, hometown, etc.- It is not tagged content, etc.*). This effectively resolves to an improved 'registration plugin'
 - Applications cannot request users give them 'contact information' including UID, email, phone; however, we do need to allow applications some way of reaching their users in a stable manner using Facebook (e.g. At a minimum spam folder in inbox, or a proxy email, or something)
 - Applications cannot request users give them 'friends' of the user (which isn't purely the user's data because it requires confirmation, etc.) ; however, (1) we allow a given app to get 'friends of this user who are also using this application' and (2) we do provide a paid 'invite' channel whereby applications can ask a user to invite more of their relationships to the application.
 - Applications cannot request information about a user's 'friends' at all, nor can they request feed, etc.
 - Our Data
 - Facebook has information about users which can be helpful to applications, and which we provide to applications when we deem appropriate, this information includes:
 - Aggregate data about the tastes, properties, etc. of a user's friends (things this user's friends like, places they live, etc.)
 - Derived data about a user / facebook's data/opinion of a user (probable location, account trust score, account age, etc.)
 - (*Data provided by third parties — information which third parties have contributed to the graph on behalf of a user*)
 - *We allow a limited number of calls for free to Any / All applications for this data (and always free calls for the admins of the app)*

- *We charge everyone on a per-call basis for this information. We have a standard rate card, but require a deal and price based on understanding of / commitment of app re: use*
- *We do not allow things which are at all competitive to 'buy' this data from us*
- Mechanics:
 - From a mechanics perspective, the APIs function much as they do today; however, if we could start over / the way platform would ideally work, we never give applications UIDs, we only give out APP-specific hashed IDs. While more complicated, this provides a significant technical hurdle for applications trying to 'suck out' the graph and makes sure that apps using our platform are committed to our distribution channels.

(B) Extending the platform to make it more valuable for engagement by anonymous writes and targeting. **The question here is what is the full opportunity / how big a deal or valuable is it for these functions to exist.**

- By default / in our base case applications can only write data 'on behalf' of users, or on their own behalf, but they can't write data 'about' users without the user's express consent via GDP, etc.
- In an 'extended' world, we would do the following:
 - Allow applications to 'target' on Facebook any user that has interacted with them in a broader set of ways (visited page, opened photo, seen post in feed, etc.) -- not underway
 - Allow applications to pass us a email address and get back a hashed-ID for a user which they can then use to communicate with / advertise to / message user -- underway in basic form
 - Allow applications to drop pixels on anything they want on the web / etc. which fire back open graph edges against the hashed-ID (visited page, took action, etc.) which the app can set as 'only me' (aka, completely private to the application) or 'me + user' (which includes the user in the privacy of the post) -- not underway
 - Allow applications to write any other data they want directly in against a user's hashed-ID for their own targeting use AND which they can allow other businesses to target against given certain terms / black list & rev-share, etc. -- not underway

(C) How to properly 'sell' information so that we are happy with the results of doing this net of the cost of change.

- How much can we make short and long-term selling data
- What is the developer thrash / cost of moving from free to paid for new or existing APIs
- We know we can't efficiently price it... how do we price it at all, etc.

From: Douglas Purdy <dmp@fb.com>
Date: Wednesday, August 29, 2012 12:47 AM
To: Mike Vernal <vernal@fb.com>
Cc: Sam Lessin <sl@fb.com>
Subject: Re: Platform Business Model Framing

working on this.

going to see the attached to frame things up.

On Aug 28, 2012, at 6:15 PM, Mike Vernal <vernal@fb.com> wrote:

Doug - as context, we're having an mteam offsite on Tue + Wed of next week to talk about three-year-plan stuff, and one of the discussions we're going to have is around the Open Graph business model.

Sam is at burning man (not sure when he gets back), and we left it a little ambiguous about who was pulling together what, so I'd like to at least get started pulling together a deck that we can use to frame the conversation. Can you ask the PMs to pull together a few slides?

Ideally, I think we want to cover:

Read APIs:

- What read APIs are free*, and cost-recovery charging for free APIs
- Premium read APIs
- Rules around using read APIs

Write APIs:

- Free write APIs (specifically, the Open Graph APIs)
- Premium distribution APIs
- Monetizing this data (sponsored stories, contextual ads on FB, contextual ads off FB)

Developer Rules:

- Anyone can build for free, have to pay a developer fee to launch
- Goes through app review process, we review Facebook integration (edges it publishes) and app assets, goes into app center
- We re-charge yearly, re-review periodically

CRM / Data Exchange Scenarios:

- "Anonymous" write APIs - the CRM/data broker scenario that Sam is passionate about
- Monetizing this data

Not sure of the best framing (and I'm a little feverish right now), but it would be good to at least start pulling this together. Sam + I spent a couple of hours talking about this on Friday evening, and he had a framework for the read API stuff and the CRM stuff he was going to write-up. Otherwise, I think you/PMs have most/all the context.

For the read stuff, I think the basic framing was:

- We move all IDs from UIDs to per-app/hashed IDs in the system
- You can always bring your data to an app (your statuses, your photos, your events, etc.)

- You can bring your friends that are also using the app to the app; ideally you can't bring non-app-user-friends to the app (you have to go through us to invite them)
- We get rid of the ability for you to take your friends' data to an app
- We have a set of premium APIs like coefficient, SI, etc. where you can pay us per-user to get this data. We set the price.
- The CRM stuff I think is the idea that others can also write this data into the system and attach a charge for other people to use this data.

I think Sam is probably the right person to write-up the CRM stuff, but the rest I think the PMs have the context on.

-mike